

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

Appellant: Yolanta Beresnevichiene ) On Appeal to the  
Patent Application No.: 10/765,719 ) Board of Appeals  
Filed: January 26, 2004 )  
For: " Data Handling Apparatus and ) Group Art Unit: 2136  
Methods" ) Examiner: Louie, O.  
 ) Date: October 24, 2008  
 )

**BRIEF ON APPEAL (Amended)**

Mail Stop Appeal Brief - Patents  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Sir:

This is an appeal from the Final Rejection (or Final Action), dated April 22, 2008, for the above identified patent application. A Notice of Appeal was filed on July 22, 2008. The original Appeal Brief was filed on September 22, 2008. This amended Appeal Brief is filed in response to the Official Action dated October 2, 2008. Please charge the Appeal Brief fee of \$510.00 to deposit account no.  
08-2025.

**REAL PARTY IN INTEREST**

The real party in interest to the present application is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

**RELATED APPEALS AND INTERFERENCES**

Appellants submit that there are no other prior and pending appeals, interferences or judicial proceedings which may be related to, directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

**STATUS OF CLAIMS**

Claims 1 and 3-41 are present in the application, are the subject of this Appeal and are reproduced in the accompanying Claims appendix. Claims 1 and 3-41 were finally rejected in the aforementioned Final Rejection. Claim 2 has been cancelled

**STATUS OF AMENDMENTS**

No claim amendments have been offered in response to the Final Rejection.

**SUMMARY OF CLAIMED SUBJECT MATTER**

Generally speaking, the present invention provides a data handling apparatus for a computer platform using an operating system executing a process, the apparatus comprising a system call monitor for detecting predetermined system calls, and means for applying a data handling policy to the system call upon a predetermined system call being detected, whereby the data handling policy is applied for all system calls involving the writing of data outside the process [p.2, ll. 23-31].

In particular, claim 1 is directed to a data handling apparatus [400; Figure 8 1; p. 27, ll. 19-33] for a computer platform [Figure 1; p. 10, ll. 16-33] using an operating system [p. 26, l. 15 - p. 27, l. 8] executing a process, the apparatus comprising a system call monitor [402; Fig. 8; p. 27, l. 19 - p. 30, l. 13] for detecting predetermined system calls and data manipulated by the process so as to modify identifiable characteristics of the data [130, Fig. 2; p. 2, ll. 23-31; p. 15, l. 10 - p. 31, l. 15 ], and means for applying a data handling policy [p. 28, l. 9 - p. 31, l. 15] upon detecting: (1) a predetermined data type based on a tag or label [130, Fig. 2; p. 2, ll. 23-31; p. 15, l. 10 - p. 22, l. 5 ] associated with the data manipulated by the process or based on the format of the data manipulated by the

process; and (2) a predetermined system call which involves the writing of data outside the process [p. 18, l. 17 - p. 20, l. 2].

Claim 22 is directed to a data handling method for a computer platform [Figure 1; p. 10, ll. 16-33] using an operating system [p. 26, l. 15 - p. 27, l. 8] executing a process, the method comprising the steps of: detecting both (i) a predetermined data type based on a tag or label [130, Fig. 2; p. 2, ll. 23-31; p. 15, l. 10 - p. 31, l. 15] associated with the data or based on the format of the data and (ii) predetermined system calls involving the writing of data outside the process [402; Fig. 8; p. 27, l. 19 - p. 30, l. 13], and applying a data handling policy [p. 28, l. 9 - p. 31, l. 15] to a system call upon both said predetermined data type and said a predetermined system call being detected, the data handling policy being applied for all system calls involving the writing of data outside the process [p. 6, ll. 1-10; p. 18, l. 17 - p. 20, l. 2].

Claim 41 is directed to a data handling apparatus for a computer platform [Figure 1; p. 10, ll. 16-33] using an operating system [p. 26, l. 15 - p. 27, l. 8] executing a process, the apparatus comprising a system call monitor [402; Fig. 8; p. 27, l. 19 - p. 30, l. 13] for detecting predetermined system calls and data handled by the process, and a policy applicator for applying a data handling policy [p. 28, l. 9 - p. 31, l. 15] to the system call upon both (i) a predetermined data type based on a tag or label [130, Fig. 2; p. 2, ll. 23-31; p. 15, l. 10 - p. 31, l. 15] associated with the data handled by the process or based on the format of the data handled by the process and (ii) a predetermined system call which involves the writing of data outside the process [p. 6, ll. 1-10; p. 18, l. 17 - p. 20, l. 2].

**GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL**

**Issue 1: Whether Claim 39 is directed to patentable under 35 U.S.C. 101?**

**Issue 2: Whether Claims 1, 4, 5, 7, 22, 25-27 and 39 -41 are patentable under 35 U.S.C. 102(e) over US Patent 6,658,571 to O'Brien et al. (hereinafter "O'Brien")?**

**Issue 3: Whether Claims 3, 6, 23, 24 & 28 are patentable under 35 U.S.C. 103(a) in view of O'Brien in view of US Patent 6,981,140 to Choo (hereinafter "Choo")?**

**Issue 4: Whether Claims 8-12, 17-21, 29, 31-33, 36 and 37 are patentable under 35 U.S.C. 103(a) in view of O'Brien in view of US Patent 5,909,688 to Yoshioka et al. (hereinafter "Yoshioka")?**

**Issue 5: Whether Claims 13-16, 30, 34, 35, and 38 are patentable under 35 U.S.C. 103(a) in view of O'Brien in view of US Patent 6,658,572 to Johnson et al. (hereinafter "Johnson")?**

## ARGUMENT

### Issue 1: Whether Claim 39 is directed to patentable under 35 U.S.C. 101?

In the final Office Action of April 22, 2008, the Examiner rejects Claim 39 under 35 U.S.C. 101 as being directed to non-statutory subject matter. Appellant respectfully disagrees.

Claim 39 is a dependent claim. It depends from claim 22 which the Examiner does not reject on this basis. Claim 39 recites:

"A computer program stored in computer readable media for controlling a computing platform to operate in accordance with claim 22."

The Examiner asserts that "'computer readable media' appears to include non-statutory subject matter" without explaining what that non-statutory subject matter might be. This is the sort of conclusory analysis that the Supreme Court stated that Examiner should not engage in when making prior art rejections, and it is submitted that the Supreme Court would take a similar view of a conclusory rejection under 35 USC 101.

The US Patent & Trademark Office was allowed many, many claims to be issued with the formulation "computer readable media" without objection. See

recently issued US patents 7,426,737 (claim 5); 7,426,727 (claims 4-6); 7,426,716 (claim 15) to identify a few of the thousands of US patents which include the formulation “computer readable media” in their claims. If the language “computer readable media” does include non-statutory subject matter (which is denied), then the USPTO should be consistent in its allowance and rejection of claims using that formulation and should start by re-examining the thousands of issued US patents using that formulation.

Furthermore, note that the Examiner wants this claim to be rewritten with a new formulation “computer readable storage media configured to control ...” in order to make it allowable. What’s to keep another Examiner from asserting that this new formulation might include non-statutory subject matter at some point in the future? And if the Applicant were to adopt that formulation in order to make this rejection disappear, then does that place in question the enforceability of all those patents issued with the formulation “computer readable media”?

The formulation “computer readable media” has been found to be statutory in that it reads on a product which is patentable under 35 USC 101 and there is nothing in the rejection to justify a different conclusion here.

**Issue 2: Whether Claims 1, 4, 5, 7, 22, 25-27 and 39-41 are patentable under 35 U.S.C. 102(e) over US Patent 6,658,571 to O’Brien et al. (hereinafter “O’Brien”)?**

Appellants submit that “[a] claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently

described, in a single prior art reference.” MPEP 2131 quoting *Verdegaal Bros. V. Union Oil Co, of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). The Examiner is also reminded that “[the] identical invention must be shown in as complete detail as is contained in the ... claim.” MPEP 2131 quoting *Richardson v. Suzuki Motor Co.*, 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). Appellants submit that the Examiner has not shown that O’Brien teaches each (or even makes obvious) and every element as set forth in the rejected claims. In particular:

**Claims 1, 4, 5, & 7**

Claim 1 recites, *inter alia*, “data manipulated by the process so as to modify identifiable characteristics of the data, and means for applying a data handling policy upon detecting:

(1) a predetermined data type based on a tag or label associated with the data manipulated by the process or based on the format of the data manipulated by the process; and ...”

The Examiner deals with this limitation on page 4 of the final rejection citing col. 4, ll. 40-48 of O’Brien. O’Brien describes a security framework for dynamically wrapping standard, commercially software applications executing in a computing system in order to limit the amount of potential damage that a successful attacker might do. O’Brien mentions “labels” and the Examiner focuses on this word apparently since claim 1 uses the same word. In O’Brien the Examiner points to a passage where O’Brien tells the reader:

“Additionally, security framework 101 can provide audit and monitoring functionality to record information regarding the extent to

which each application 107 accesses computing resources 106. Also, security framework 101 can be used to implement more sophisticated label-based policies, such as a multilevel secure (MLS) policy or a type enforcement policy. In one embodiment, security framework 101 creates and maintains a list of labels, each label corresponding to a computing resource 106. Because of the flexibility of security framework 101, these label-based policies could be quickly changed, if needed, to adapt to new operating requirements.” (O’Brien, col 4, ll. 37-48)

This passage seems to discuss how the O’Brien disclosure are be modified or utilized by using “more sophisticated label-based policies” which the reader is left to his or her devices to figure out just what the applicant in O’Brien is trying to say or how this added feature is to be implemented in the basic O’Brien disclosure.

In any event, the reader is told that the labels correspond to a computing resource 106, which O’Brien tells the reader earlier on are such things as “memory, files, network sockets and processes” (O’Brien col. 3, ll. 8-9). Note that the reader is told that O’Brien’s labels allegedly correspond to “a” computing resource. Which of one of the “memory, files, network sockets and processes” that the “a” computing resource that O’Brien’s label might correspond to is left unstated.

Importantly, the claim limitation in question recites “a predetermined data type based on a tag or label associated with the data manipulated by the process or based on the format of the data manipulated by the process”. It is submitted that those skilled in the art do not think of “the data manipulated by the process” or “ the format of the data manipulated by the process” as being “a computing

resource" and therefore the attempt at equating the two is quite unreasonable. Based on O'Brien, a computing resource can well be memory, but there is no reason to associate the term "computing resource" with data stored in memory based on O'Brien's disclosure. In short, O'Brien does not lead one to associate labels "with the data manipulated by the process or based on the format of the data manipulated by the process" as required by claim 1.

Recall that "[the] identical invention must be shown in as complete detail as is contained in the ... claim" and the disclosure in O'Brien is anything but clear. O'Brien certainly does not lead unambiguously to claim 1. Indeed, it is submitted that O'Brien is not an enabling disclosure when it comes to the issue of his labels. How those labels might be used and what sort of computer resource they might be associated with is highly speculative and left to the reader to try to figure out. The Examiner just speculates that O'Brien's disclosure anticipates claim 1 without bothering to read the specific language of claim 1 on O'Brien.

Claim includes other limitations such as "a system call monitor for detecting predetermined system calls and data manipulated by the process so as to modify identifiable characteristics of the data".

The Examiner asserts that this limitation is anticipated by the following disclosure of O'Brien:

"Security manager 111 provides an interface for communicating with security master 103, thereby allowing user 113 to configure and control security modules 105 from user space 104. Management functions available to user 113 include: the ability to list a set of rules that are being

enforced by each security module 105, the ability to load a new set of rules for a particular security module 105, and the ability to log and view activity within security framework 101" (O'Brien col. 4, ll. 5-14).

This is very interesting, but this disclosure in O'Brien does not anticipate the above-quoted limitation regarding the call monitor. The Examiner's discussion about wrapping a web browser (see page 4, line 16 of the final rejection) is more on point, but the Examiner basically assumes that the limitation a system call monitor for detecting predetermined system calls and data manipulated by the process so as to modify identifiable characteristics of the data" is met without explaining how it is met by O'Brien. The Examiner's analysis seems to concentrate on the system calls portion of claim 1 and seemingly ignores the requirement of claim 1 for "a system call monitor for detecting ... data manipulated by the process so as to modify identifiable characteristics of the data."

But even if it is assumed that this limitation is met by O'Brien, O'Brien clearly fails to met the limitation regarding "a predetermined data type based on a tag or label associated with the data manipulated by the process or based on the format of the data manipulated by the process" as discussed above. Note also the lack of a connection between O'Brien's discussion of "more sophisticated label-based policies" and his discussion of wrapping a web browser. The two are apparently unrelated, yet the Examiner acts as if they are for the purposes of this rejection. Note that claim 1 recites:

"means for applying a data handling policy upon detecting:

(1) a predetermined data type based on a tag or label associated with the data manipulated by the process or based on the format of the data manipulated by the process; and

(2) a predetermined system call which involves the writing of data outside the process."

Both detections of subparagraphs (1) and (2) must occur to trigger the means for applying a data handling policy. No such teaching occurs in O'Brien. Claim 1 (and claims dependent thereon) is patentable over O'Brien.

Claim 22, 25-27 & 39-41

The Examiner cites basically the same passages in column 4 of O'Brien as his basis for rejecting claim 22 (and claims dependent thereon) as he cited in making his rejection of claim 1. See pages 6 and 7 of the final rejection. Claim 22 recites, *inter alia*, "detecting both (i) a predetermined data type based on a tag or label associated with the data or based on the format of the data and (ii) predetermined system calls involving the writing of data outside the process, and applying a data handling policy to a system call upon both said predetermined data type and said a predetermined system call being detected, the data handling policy being applied for all system calls involving the writing of data outside the process."

The Examiner again focuses on O'Brien's discussion of "sophisticated label-based policies" found at col. 4, ll. 37-48. As mentioned above, relative to the

rejection of claim 1, the reader of the O'Brien patent is left to his or her devices to figure out just what O'Brien is trying to say or how the "sophisticated label-based policies" feature is to be implemented into the basic O'Brien disclosure.

And as previously mentioned, the reader must try to figure out which of one of the "memory, files, network sockets and processes" that comprises computer resources O'Brien is referring to when he makes reference to "a computing resource" that O'Brien's label might correspond to.

Importantly, the claim limitation of claim 22 in question recites "a predetermined data type based on a tag or label associated with the data or based on the format of the data". It is submitted that those skilled in the art do not think of "the data" or "the format of the data" as being "a computing resource" and therefore the attempt at equating the two is quite unreasonable.

Recall again that "[the] identical invention must be shown in as complete detail as is contained in the ... claim" and the disclosure in O'Brien is anything but clear. O'Brien certainly does not lead unambiguously to claim 22. Indeed, it is submitted that O'Brien is not an enabling disclosure when it comes to the issue of his labels. How those labels might be used and what sort of computer resource they might be associated with is highly speculative and left to the reader to try to figure out. The Examiner just speculates that O'Brien's disclosure anticipates claim 22 without bothering to read the specific language of claim 22 on O'Brien.

Claim 22 also calls for "detecting both" the subject matter of subparagraph (i) of claim 22 and the subject matter of subparagraph (ii) of claim 22. The Examiner tries to read subject matter of subparagraph (i) of claim 22 upon O'Brien's confusing discussion of labels and the subject matter of subparagraph

(ii) of claim 22 on O'Brien's wrapping a web browser. But there is no connection between O'Brien's wrapped web browser and his discussion of labels which leads unambiguously to "detecting both" the subject matter of subparagraph (i) of claim 22 and the subject matter of subparagraph (ii) of claim 22 before "applying a data handling policy to a system call upon both said predetermined data type and said a predetermined system call being detected" as specifically required by claim 22.

Claim 22 (and claims dependent thereon) is patentable over O'Brien.

#### Claim 41

The Examiner cites basically the same passages in column 4 of O'Brien as his basis for rejecting claim 41 as he cited in making his rejection of claims 1 and 22. See pages 9 and 10 of the final rejection. Claim 41 recites, *inter alia*, "a policy applicator for applying a data handling policy to the system call upon both (i) a predetermined data type based on a tag or label associated with the data handled by the process or based on the format of the data handled by the process and (ii) a predetermined system call which involves the writing of data outside the process."

The Examiner focuses once more on O'Brien's discussion of "sophisticated label-based policies" found at col. 4, ll. 37-48. As mentioned above, relative to the rejection of claims 1 and 22, the reader of the O'Brien patent is left to his or her devices to figure out just what O'Brien is trying to say or how the "sophisticated label-based policies" feature is to be implemented into the basic O'Brien disclosure.

And as previously mentioned, the reader must try to figure out which of one of the “memory, files, network sockets and processes” that comprises computer resources O’Brien is referring to when he makes reference to “a computing resource” that his labels might correspond to.

Importantly, the claim limitation of claim 41 in question recites “a predetermined data type based on a tag or label associated with the data handled by the process or based on the format of the data handled by the process”. It is submitted that those skilled in the art do not think of “the data” or “the format of the data” as being “a computing resource” and therefore the attempt at equating the two is quite unreasonable.

Recall once more that “[the] identical invention must be shown in as complete detail as is contained in the ... claim” and the disclosure in O’Brien is anything but clear. O’Brien certainly does not lead unambiguously to claim 41. Indeed, it is submitted that O’Brien is not an enabling disclosure when it comes to the issue of his labels. How those labels might be used and what sort of computer resource they might be associated with is highly speculative and left to the reader to try to figure out. The Examiner just speculates that O’Brien’s disclosure anticipates claim 41 without bothering to read the specific language of claim 41 on O’Brien.

Claim 41 also calls for “a policy applicator for applying a data handling policy to the system call upon both” the subject matter of subparagraph (i) of claim 41 and the subject matter of subparagraph (ii) of claim 41. The Examiner tries to read subject matter of subparagraph (i) of claim 41 upon O’Brien’s confusing discussion of labels and the subject matter of subparagraph (ii) of

claim 41 on O'Brien's wrapping a web browser. But there is no connection between O'Brien's wrapped web browser and his discussion of labels which leads unambiguously to "a policy applicator for applying a data handling policy to the system call upon both" the subject matter of subparagraph (i) of claim 41 and the subject matter of subparagraph (ii) of claim 41 occurring as specifically required by claim 41.

Claim 41 is patentable over O'Brien.

**Issue 3: Whether Claims 3, 6, 23, 24 & 28 are patentable under 35 U.S.C. 103(a) in view of O'Brien in view of US Patent 6,981,140 to Choo (hereinafter "Choo")?**

Claims 6 and 28

Claim 6 recites a "data handling apparatus according to claim 5, in which the policy interpreter is configured to use the intended destination of the data as a factor in determining the policy for the data." Claim 28 recites a "data handling method according to claim 26, in which the intended destination of the data is used as a factor in determining the policy for the data." The Examiner cites col. 12, ll. 54-59 of Choo which mentions that for incoming data packets from a remote host that the packets are inspected to see if they are encrypted (so that decryption is needed). So the Examiner reasons that Choo teaches "policy enforcement / access control based on where the packets come from" (see page 10, last line of the final rejection) and then seems to suggest that this is somehow

analogous to the claim language quoted above. There are at least two problems with this analysis.

First, testing for a descriptor showing a packet to be encrypted is not the same thing as testing “where the packets come from”.

Second, assuming for the moment that the Examiner is correct on the “where the packets come from” issue, the Examiner’s justification for turning this around to examining where the packets are going (to try to meet the limitation “policy interpreter is configured to use the intended destination of the data as a factor in determining the policy for the data”) just does not exist.

35 U.S.C. § 103 “forbids issuance of a patent when ‘the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains.’” *KSR Int’l Co. v. Teleflex Inc.*, 127 S.Ct. 1727, 1734 (2007). The Supreme Court stated that obvious analysis “should be made explicit.” The analysis provided to support this rejection of claims 6 and 28 is anything but explicit.

The rejection of the claims 6 and 28 fails to comply with Supreme Court’s mandate in KSR.

**Issue 4: Whether Claims 8-12, 17-21, 29, 31-33, 36 and 37 are patentable under 35 U.S.C. 103(a) in view of O’Brien in view of US Patent 5,909,688 to Yoshioka et al. (hereinafter “Yoshioka”)?**

Claims 8-12, 17-21, 29, 31-33, 36 and 37

The Examiner provides a rather confusing discussion of what Yoshioka teaches on pages 13 and 14 of the final rejection. The Examiner's rationale for combining O'Brien and Yoshioka is stated on page 14 of the final rejection. This discussion is merely conclusory.

The Examiner mentions on the top of page 14 "(the data management unit arranged to) regulate operating system operations involving the data ..." and then delves into a discussion of Figure 13 of Yoshioka.

But the Examiner gives short shrift the requirements of claim 8 that "the data management unit arranged to associate data management information with data input to a process, and regulate operating system operations involving the data according to the data management information." Similar limitations can be found in claim 29. Somehow the requirement for "data input to a process" gets associated with data corresponding to each record in Yoshioka's database. Why are these supposedly the same thing?

And why add a database to O'Brien's security framework where he dynamically wraps ordinary applications? The stated justification combining these two disclosures is "for the purposes of associating and tracking data processed in an operating system" (see the last two lines on page 14 of the final rejection). This is a mere conclusory statement. What data needs to be associated and tracked in O'Brien and why do that?

As previously mentioned, 35 U.S.C. § 103 "forbids issuance of a patent when 'the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to

which said subject matter pertains.” *KSR Int'l Co. v. Teleflex Inc.*, 127 S.Ct. 1727, 1734 (2007). The Supreme Court stated that obvious analysis “should be made explicit.” The analysis provided to support this combining of O’Brien and Yoshioka is anything but explicit and the analysis provided by the Examiner in the final rejection is very confusing. The Applicant is entitled to a well reasoned rationale for rejecting claims that the Applicant can then consider how best to respond to. The Applicant is not given a fair opportunity to respond to a rejection when the underlying rationale is both confusing and conclusory.

The Examiner has not provided a proper rationale for combining O’Brien and Yoshioka and therefore the rejections based on same should be overturned.

**Issue 5: Whether Claims 13-16, 30, 34, 35, and 38 are patentable under 35 U.S.C. 103(a) in view of O'Brien in view of US Patent 6,658,572 to Johnson et al. (hereinafter “Johnson”)?**

Claims 13-16

Claim 13 recites a “data handling apparatus according to claim 8, in which the data management unit comprises part of an operating system kernel space.”

The Examiner admits on page 13 of the final rejection that O’Brien does not meet the “data management unit” limitation. So this rejection is fatally defective since the Examiner does not suggest that Johnson makes up for this deficiency of O’Brien.

When rejecting claim 8, upon which claims 13-16 depend (directly or indirectly), the Examiner also relied upon Yoshioka, which he does not rely upon for this rejection. And if the Applicant is supposed to assume that the Examiner really intended to cite a combination of O'Brien, Yoshioka and Johnson in rejecting claims 13-16, then the final rejection hardly meets the KSR test requiring that the rejection be made explicit.

Moreover, this rejection would then need to include a suitable justification for taking Yoshioka's database code and stuffing that into an operating system kernel. Where is that analysis? It is missing! Does the Examiner have some prior art document suggesting that a database such as Yoshioka's should be stuff into an operating system kernel? If so, let him cite it!

#### Claims 30, 34, 35 & 38

When rejecting claim 29, upon which claims 30, 34, 35 and 38 depend (directly or indirectly), the Examiner also relied upon Yoshioka, which he does not rely upon for this rejection<sup>1</sup>. And if the Applicant is supposed to assume that the Examiner really intended to cite a combination of O'Brien, Yoshioka and Johnson in rejecting claims 30, 34, 35 and 38, then the final rejection hardly meets the KSR test requiring that the rejection be made explicit.

Moreover, this rejection would then need to include a suitable justification for taking Yoshioka's database code and combining that with the teachings of

---

<sup>1</sup> The Examiner mentions Yoshioka at page 29 line 4 (and elsewhere), but the reference(s) is (are) apparently to Johnson and not Yoshioka.

Johnson. Where is that analysis? It is missing! As previously mentioned, Yoshioka fails to deal with the limitation of claim 29 requiring “associating data management information with data input to a process” and thus how is the limitation of claim 30 (for example) that “supervisor code administers the method by controlling the process at run time” supposed to be dealt with? What process in O’Brien/Yoshioka/Johnson is supposed to read on the rejected claim language?

Again, 35 U.S.C. § 103 “forbids issuance of a patent when ‘the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains.’” *KSR Int'l Co. v. Teleflex Inc.*, 127 S.Ct. 1727, 1734 (2007). The Supreme Court stated that obvious analysis “should be made explicit.” The analysis provided to support this rejection of claims 30, 34, 35 and 38 is anything but explicit.

### Conclusion

For the extensive reasons advanced above, Appellants respectfully contend that each claim is patentable over the cited prior art and that claim 39 is statutory. Therefore, reversal of all rejections is courteously solicited.

The Commissioner is authorized to charge any additional fees which may be required or credit overpayment to deposit account no. 08-2025. In particular, if this Appeal Brief is not timely filed, the Commissioner is authorized to treat this response as including a petition to extend the time period pursuant to 37 CFR 1.136(a) requesting an extension of time of the number of months necessary to make this response timely filed and the petition fee due in connection therewith may be charged to deposit account no. 08-2025.

Respectfully submitted,

I hereby certify that this correspondence is being electronically filed with the United States Patent and Trademark Office on

24 October 2008

(Date of Transmission)

Lonnie Louie

(Name of Person Transmitting)

/Lonnie Louie/

(Signature)

24 October 2008

(Date))

/Richard P. Berg 28145/

Richard P. Berg  
Attorney for the Applicant  
Reg. No. 28,145  
LADAS & PARRY  
5670 Wilshire Boulevard,  
Suite 2100  
Los Angeles, California 90036  
(323) 934-2300 voice  
(323) 934-0202 facsimile

Encls:

Claims Appendix;  
Evidence Appendix;  
Related Proceedings Appendix.

1. A data handling apparatus for a computer platform using an operating system executing a process, the apparatus comprising a system call monitor for detecting predetermined system calls and data manipulated by the process so as to modify identifiable characteristics of the data, and means for applying a data handling policy upon detecting:

(1) a predetermined data type based on a tag or label associated with the data manipulated by the process or based on the format of the data manipulated by the process; and

(2) a predetermined system call which involves the writing of data outside the process.

3. A data handling apparatus according to claim 6, in which the policy interpreter in its application of the policy automatically encrypts the at least some of the data.

4. A data handling apparatus according to claim 1, in which predetermined system calls are those involving the transmission of data externally of the computing platform.

5. A data handling apparatus according to claim 1, in which the means for applying a data handling policy comprises a tag determiner for determining any security tags associated with the data manipulated by the process or based on the format of the data manipulated by the process handled by the system call, and a policy interpreter for determining a policy according to any such security tags and for applying the policy.

6. A data handling apparatus according to claim 5, in which the policy interpreter is configured to use the intended destination of the data as a factor in determining the policy for the data.

7. A data handling apparatus according to claim 5, in which the policy interpreter comprises a policy database including tag policies and a policy reconciler for generating a composite policy from the tag policies relevant to the data.
8. A data handling apparatus according to claim 1, in which the computing platform comprises a data management unit, the data management unit arranged to associate data management information with data input to a process, and regulate operating system operations involving the data according to the data management information.
9. A data handling apparatus according to claim 8, in which the computing platform further comprises a memory space, and is arranged to load the process into the memory space and run the process under the control of the data management unit.
10. A data handling apparatus according to claim 8, in which the data management information is associated with at least one data sub-unit as data is input to a process from a data unit comprising a plurality of sub-units.
11. A data handling apparatus according to claim 8, in which data management information is associated with each independently addressable data unit.
12. A data handling apparatus according to claim 8, in which the data management unit comprises part of an operating system kernel space.

13. A data handling apparatus according to claim 12, in which the operating system kernel space comprises a tagging driver arranged to control loading of a supervisor code into the memory space with the process.
14. A data handling apparatus according to claim 13, in which the supervisor code controls the process at run time to administer the operating system data management unit.
15. A data handling apparatus according to claim 14, in which the supervisor code is arranged to analyse instructions of the process to identify operations involving the data, and, provide instructions relating to the data management information with, the operations involving the data.
16. A data handling apparatus according to claim 13, in which the memory space further comprises a data management information area under control of the supervisor code arranged to store the data management information.
17. A data handling apparatus according to claim 8, in which the data management unit comprises a data filter to identify data management information associated with data that is to be read into the memory space.
18. A data handling apparatus according to claim 8, in which the data management unit further comprises a tag management module arranged to allow a user to specify data management information to be associated with data.

19. A data handling apparatus according to claim 8, in which the data management unit comprises a tag propagation module arranged to maintain an association with the data that has been read into the process and the data management information associated therewith.

20. A data handling apparatus according to claim 19, in which the tag propagation module is arranged to maintain an association between an output of operations carried out within the process and the data management information associated with the data involved in the operations.

21. A data handling apparatus according to claim 19, in which the tag propagation module comprises state machine automata arranged to maintain an association between an output of operations carried out within the process and the data management information associated with the data involved in the operations.

22. A data handling method for a computer platform using an operating system executing a process, the method comprising the steps of:

detecting both (i) a predetermined data type based on a tag or label associated with the data or based on the format of the data and (ii) predetermined system calls involving the writing of data outside the process, and applying a data handling policy to a system call upon both said predetermined data type and said a predetermined system call being detected, the data handling policy being applied for all system calls involving the writing of data outside the process.

23. A data handling method according to claim 22, in which the policy is to require the encryption of at least some of the data.
24. A data handling method according to claim 23, in which in its application of the policy at least some of the data is automatically encrypted.
25. A data handling method according to claim 22, in which predetermined system calls are those involving the transmission of data externally of the computing platform.
26. A data handling method according to claim 22, in which the method includes the steps of: determining any security tags associated with data handled by the system call, determining a policy according to any such tags and applying the policy.
27. A data handling method according to claim 26, in which a composite policy is generated from the tag policies relevant to the data.
28. A data handling method according to claim 26, in which the intended destination of the data is used as a factor in determining the policy for the data.
29. A data handling method according to claim 22, in which the method further comprises the steps of: (a) associating data management information with data input to a process; and (b) regulating operating system operations involving the data according to the data management information.

30. A data handling method according to claim 29, in which supervisor code administers the method by controlling the process at run time.
31. A data handling method according to claim 29, in which the step (a) comprises associating data management information with data as the data is read into a memory space.
32. A data handling method according to claim 29, in which the step (a) comprises associating data management information with at least one data sub-unit as data is read into a memory space from a data unit comprising a plurality of data sub-units.
33. A data handling method according to claim 29, in which the step (a) comprises associating data management information with each independently addressable data unit that is read into the memory space.
34. A data handling method according to claim 29, in which the data management information is written to a data management memory space under control of the supervisor code.
35. A data handling method according to claim 34, in which the supervisor code comprises state machine automata arranged to control the writing of data management information to the data management memory space.
36. A data handling method according to claim 29, in which the step (b) comprises sub-steps (b1) identifying an operation involving the data; (b2) if the operation involves the data and is carried out within the process, maintaining an association

between an output of the operation and the data management information; and (b3) if the operation involving the data includes a write operation to a location external to the process, selectively performing the operation dependent on the data management information.

37. A data handling method according to claim 36, in which, the step (b1) comprises: analysing process instructions to identify operations involving the data; and, providing instructions relating to the data management information with the operations involving the data.

38. A data handling method according to claim 29, in which the process instructions are analysed as blocks, each block defined by operations up to a terminating condition.

39. A computer program stored in computer readable media for controlling a computing platform to operate in accordance with claim 22.

40. A computer platform configured to operate according to claim 22.

41. A data handling apparatus for a computer platform using an operating system executing a process, the apparatus comprising a system call monitor for detecting predetermined system calls and data handled by the process, and a policy applicator for applying a data handling policy to the system call upon both (i) a predetermined data type based on a tag or label associated with the data handled by the process or based on the format of

the data handled by the process and (ii) a predetermined system call which involves the writing of data outside the process.

No evidence is being submitted

No copies of decisions rendered in related proceedings are being submitted.